

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

AN INTERFACE AND RELATED METHODS

FOR

**DYNAMICALLY GENERATING A FILTER GRAPH IN A
DEVELOPMENT SYSTEM**

Inventor(s):

Daniel J. Miller

Eric H. Rudolph

ATTORNEY'S DOCKET NO. MS1-634US

1 **TECHNICAL FIELD**

2
3 This invention generally relates to processing media content and, more
4 particularly, to an interface and related methods for dynamically generating a filter
5 graph in a development system.

6 **BACKGROUND**

7
8 Recent advances in computing power and related technology have fostered
9 the development of a new generation of powerful software applications. Gaming
10 applications, communications applications, and multimedia applications have
11 particularly benefited from increased processing power and clocking speeds.
12 Indeed, once the province of dedicated, specialty workstations, many personal
13 computing systems now have the capacity to receive, process and render
14 multimedia objects (e.g., audio and video content). While the ability to display
15 (receive, process and render) multimedia content has been around for a while, the
16 ability for a standard computing system to support true multimedia editing
17 applications is relatively new.

18 In an effort to satisfy this need, Microsoft Corporation introduced an
19 innovative development system supporting advanced user-defined multimedia
20 editing functions. An example of this architecture is presented in US Patent No.
21 5,913, 038 issued to Griffiths and commonly owned by the assignee of the present
22 invention, the disclosure of which is expressly incorporated herein by reference.

23 In the '038 patent, Griffiths introduced the an application program interface
24 which, when exposed to higher-level development applications, enable a user to
25 graphically construct a multimedia processing project by piecing together a

1 collection of "filters" exposed by the interface. The interface described therein is
2 referred to as a filter graph manager. The filter graph manager controls the data
3 structure of the filter graph and the way data moves through the filter graph. The
4 filter graph manager provides a set of component object model (COM) interfaces
5 for communication between a filter graph and its application. Filters of a filter
6 graph architecture are preferably implemented as COM objects, each
7 implementing one or more interfaces, each of which contains a predefined set of
8 functions, called methods. Methods are called by an application program or other
9 component objects in order to communicate with the object exposing the interface.
10 The application program can also call methods or interfaces exposed by the filter
11 graph manager object.

12 Filter graphs work with data representing a variety of media (or non-media)
13 data types, each type characterized by a data stream that is processed by the filter
14 components comprising the filter graph. A filter positioned closer to the source of
15 the data is referred to as an upstream filter, while those further down the
16 processing chain is referred to as a downstream filter. For each data stream that
17 the filter handles it exposes at least one virtual pin (i.e., distinguished from a
18 physical pin such as one might find on an integrated circuit). A virtual pin can be
19 implemented as a COM object that represents a point of connection for a
20 unidirectional data stream on a filter. Input pins represent inputs and accept data
21 into the filter, while output pins represent outputs and provide data to other filters.
22 Each of the filters include at least one memory buffer, wherein communication of
23 the media stream between filters is accomplished by a series of "copy" operations
24 from one filter to another.

As introduced in Griffiths, a filter graph has three different types of filters: source filters, transform filters, and rendering filters. A source filter is used to load data from some source; a transform filter processes and passes data; and a rendering filter renders data to a hardware device or other locations (e.g., saved to a file, etc.). An example of a filter graph for a simplistic media rendering process is presented with reference to Fig. 1.

Fig. 1 graphically illustrates an example filter graph for rendering media content. As shown, the filter graph 100 is comprised of a plurality of filters 102-114, which read, process (transform) and render media content from a selected source file. As shown, the filter graph includes each of the types of filters described above, interconnected in a linear fashion.

Products utilizing the filter graph have been well received in the market as it has opened the door to multimedia editing using otherwise standard computing systems. It is to be appreciated, however, that the construction and implementation of the filter graphs are computationally intensive and expensive in terms of memory usage. Even the most simple of filter graphs requires and abundance of memory to facilitate the copy operations required to move data between filters. Thus, complex filter graphs can become unwieldy, due in part to the linear nature of conventional development system architecture. Moreover, it is to be appreciated that the filter graphs themselves consume memory resources, thereby compounding the issue introduced above.

Thus, what is required is a filter graph architecture which reduces the computational and memory resources required to support even the most complex of multimedia projects. Indeed, what is required is a development interface and related methods that dynamically generates a filter graph during project execution,

1 thereby improving the perceived performance of the development system. Just
2 such a solution is disclosed below.

3 4 **SUMMARY**

5 This invention concerns a system and related interfaces supporting the
6 processing of media content. In accordance with one aspect of the present
7 embodiment, a system system is presented including a plurality of sources, and an
8 interface, selectively coupled to the plurality of sources, to generate and
9 implement a development project of processing chains, wherein the interface
10 dynamically loads a processing chain for each of the plurality of media sources at
11 a point during the execution of the project when the chain is required, and wherein
12 the interface is configured to unload at least a subset of the chains when they are
13 not required. By dynamically modifying the development project to include only
14 those chains necessary to support the execution of the project, the interface
15 reduces the computational and memory load placed on the host system.

16 17 **BRIEF DESCRIPTION OF THE DRAWINGS**

18 The same reference numbers are used throughout the figures to reference
19 like components and features.

20 Fig. 1 is a graphical representation of a conventional filter graph
21 representing a user-defined development project.

22 Fig. 2 is a block diagram of a computing system incorporating the teachings
23 of the described embodiment.

24 Fig. 3 is a block diagram of an example software architecture incorporating
25 the teachings of the described embodiment.

Fig. 4 is a graphical illustration of an example software-enabled matrix switch, according to an exemplary embodiment.

Fig. 5 is a graphical representation of a data structure comprising a programming grid to selectively couple one or more of a scalable plurality of input pins to a scalable plurality of output pins of the matrix switch filter, in accordance with one aspect of the described embodiment.

Fig. 6 is a graphical illustration denoting shared buffer memory between filters, according to one aspect of the described embodiment.

Fig. 7 is a flow chart of an example method for generating a filter graph, in accordance with one aspect of the described embodiment.

Fig. 8 is a flow chart of an example method for negotiating buffer requirements between at least two adjacent filters, according to one aspect of the described embodiment.

Fig. 9 graphically illustrates an overview of a process that takes a user-defined editing project and composites a data structure that can be used to program the matrix switch.

Fig. 10 graphically illustrates the project of Fig. 9 in greater detail.

Fig. 11 shows an exemplary matrix switch dynamically generated in support of the project developed in Figs. 9 and 10, according to one described embodiment.

Fig. 12 illustrates a graphic representation of an exemplary data structure that represents the project of Fig. 10, according to one described embodiment.

Figs. 13-18 graphically illustrate various states of a matrix switch programming grid at select points in processing the project of Figs. 9 and 10 through the matrix switch, in accordance with one described embodiment.

Fig. 19 is a flow chart of an example method for processing media content, in accordance with one described embodiment.

Fig. 20 illustrates an example project with a transition and an effect, in accordance with one described embodiment.

Fig. 21 shows an exemplary data structure in the form of a hierarchical tree that represents the project of Fig. 20.

Figs. 22 and 23 graphically illustrate an example matrix switch programming grid associated with the project of Fig. 20 at select points in time, according to one described embodiment.

Fig. 24 shows an example matrix switch dynamically generated and configured as the grid of Figs. 22 and 23 was being processed, in accordance with one described embodiment.

Fig. 25 shows an exemplary project in accordance with one described embodiment.

Fig. 26 graphically illustrates an example audio editing project, according to one described embodiment.

Fig. 27 depicts an example matrix switch programming grid associated with the project of Fig. 26.

Fig. 28 shows an example matrix switch dynamically generated and configured in accordance with the programming grid of Fig. 27 to perform the project of Fig. 26, according to one described embodiment.

Fig. 29 illustrates an exemplary media processing project incorporating another media processing project as a composite, according to yet another described embodiment.

Fig. 30 graphically illustrates an example data structure in the form of a hierarchical tree structure that represents the project of Fig. 29.

Figs 31-36 graphically illustrate various matrix switch programming grid states at select points in generating and configuring the matrix switch to implement the media processing of Fig. 29.

Fig. 38 illustrates an example matrix switch suitable for use in the media processing project of Fig. 29, according to one described embodiment.

Fig. 38a graphically illustrates an example data structure in the form of a hierarchical tree structure that represents a project that is useful in understanding composites in accordance with the described embodiments.

Fig. 39 is a flow diagram that describes steps in a method in accordance with one described embodiment.

Fig. 40 is a flow chart of an example method for processing media content, in accordance with another embodiment of the present invention.

Fig. 41 is a flow chart of an example method for dynamically generating a filter graph during execution of a development project, according to one aspect of the present invention.

Fig. 42 illustrates an example data structure utilized to manage dynamic graph building, according to one embodiment.

Fig. 43 graphically illustrates a filter graph during dynamic graph building, according to one example implementation.

Fig. 44 graphically illustrates a filter graph with chain dependencies during dynamic graph building, according to one embodiment of the present invention.

DETAILED DESCRIPTION

Related Applications

This application is related to the following commonly-filed U.S. Patent Applications, all of which are commonly assigned to Microsoft Corp., the disclosures of which are incorporated by reference herein:

- Application Serial No. _____, entitled “An Interface and Related Methods for Reducing Source Accesses in a Development System”, naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-643US;
- Application Serial No. _____, entitled “A System and Related Methods for Reducing Source Filter Invocation in a Development Project”, naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-631US;
- Application Serial No. _____, entitled “A System and Related Methods for Reducing Memory Requirements of a Media Processing System”, naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-632US;
- Application Serial No. _____, entitled “A System and Related Methods for Reducing the Instances of Source Files in a Filter Graph”, naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-633US;
- Application Serial No. _____, entitled “A System and Related Methods for Processing Audio Content in a Filter Graph”, naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-639US;
- Application Serial No. _____, entitled “A System and Methods for Generating an Managing Filter Strings in a Filter Graph”, naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-642US;
- Application Serial No. _____, entitled “Methods and Systems for Processing Media Content”, naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-640US;
- Application Serial No. _____, entitled “Systems for Managing Multiple Inputs and Methods and Systems for Processing Media Content ”, naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-635US;

- Application Serial No. _____, entitled “Methods and Systems for Implementing Dynamic Properties on Objects that Support Only Static Properties”, naming Daniel J. Miller and David Maymudes as inventors, and bearing attorney docket number MS1-638US;
- Application Serial No. _____, entitled “Methods and Systems for Efficiently Processing Compressed and Uncompressed Media Content”, naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-630US;
- Application Serial No. _____, entitled “Methods and Systems for Effecting Video Transitions Represented By Bitmaps”, naming Daniel J. Miller and David Maymudes as inventors, and bearing attorney docket number MS1-637US;
- Application Serial No. _____, entitled “Methods and Systems for Mixing Digital Audio Signals”, naming Eric H. Rudolph as inventor, and bearing attorney docket number MS1-636US; and
- Application Serial No. _____, entitled “Methods and Systems for Processing Multi-media Editing Projects”, naming Eric H. Rudolph as inventor, and bearing attorney docket number MS1-641US.

Various described embodiments concern an application program interface associated with a development system. According to one example implementation, the interface is exposed to a media processing application to enable a user to dynamically generate complex media processing tasks, e.g., editing projects. In the discussion herein, aspects of the invention are developed within the general context of computer-executable instructions, such as program modules, being executed by one or more conventional computers. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, personal digital assistants, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe

1 computers, and the like. In a distributed computer environment, program modules
2 may be located in both local and remote memory storage devices. It is noted,
3 however, that modification to the architecture and methods described herein may
4 well be made without deviating from spirit and scope of the present invention.
5 Moreover, although developed within the context of a media processing system
6 paradigm, those skilled in the art will appreciate, from the discussion to follow,
7 that the application program interface may well be applied to other development
8 system implementations. Thus, the media processing system described below is
9 but one illustrative implementation of a broader inventive concept.

11 **Example System Architecture**

12 **Fig. 2** illustrates an example of a suitable computing environment 200 on
13 which the system and related methods for processing media content may be
14 implemented.

15 It is to be appreciated that computing environment 200 is only one example
16 of a suitable computing environment and is not intended to suggest any limitation
17 as to the scope of use or functionality of the media processing system. Neither
18 should the computing environment 200 be interpreted as having any dependency
19 or requirement relating to any one or combination of components illustrated in the
20 exemplary computing environment 200.

21 The media processing system is operational with numerous other general
22 purpose or special purpose computing system environments or configurations.
23 Examples of well known computing systems, environments, and/or configurations
24 that may be suitable for use with the media processing system include, but are not
25 limited to, personal computers, server computers, thin clients, thick clients, hand-

1 held or laptop devices, multiprocessor systems, microprocessor-based systems, set
2 top boxes, programmable consumer electronics, network PCs, minicomputers,
3 mainframe computers, distributed computing environments that include any of the
4 above systems or devices, and the like.

5 In certain implementations, the system and related methods for processing
6 media content may well be described in the general context of computer-
7 executable instructions, such as program modules, being executed by a computer.
8 Generally, program modules include routines, programs, objects, components,
9 data structures, etc. that perform particular tasks or implement particular abstract
10 data types. The media processing system may also be practiced in distributed
11 computing environments where tasks are performed by remote processing devices
12 that are linked through a communications network. In a distributed computing
13 environment, program modules may be located in both local and remote computer
14 storage media including memory storage devices.

15 In accordance with the illustrated example embodiment of Fig. 2 computing
16 system 200 is shown comprising one or more processors or processing units 202, a
17 system memory 204, and a bus 206 that couples various system components
18 including the system memory 204 to the processor 202.

19 Bus 206 is intended to represent one or more of any of several types of bus
20 structures, including a memory bus or memory controller, a peripheral bus, an
21 accelerated graphics port, and a processor or local bus using any of a variety of
22 bus architectures. By way of example, and not limitation, such architectures
23 include Industry Standard Architecture (ISA) bus, Micro Channel Architecture
24 (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association
25

1 (VESA) local bus, and Peripheral Component Interconnects (PCI) buss also
2 known as Mezzanine bus.

3 Computer 200 typically includes a variety of computer readable media.
4 Such media may be any available media that is locally and/or remotely accessible
5 by computer 200, and it includes both volatile and non-volatile media, removable
6 and non-removable media.

7 In Fig. 2, the system memory 204 includes computer readable media in the
8 form of volatile, such as random access memory (RAM) 210, and/or non-volatile
9 memory, such as read only memory (ROM) 208. A basic input/output system
10 (BIOS) 212, containing the basic routines that help to transfer information
11 between elements within computer 200, such as during start-up, is stored in ROM
12 208. RAM 210 typically contains data and/or program modules that are
13 immediately accessible to and/or presently be operated on by processing unit(s)
14 202.

15 Computer 200 may further include other removable/non-removable,
16 volatile/non-volatile computer storage media. By way of example only, Fig. 2
17 illustrates a hard disk drive 228 for reading from and writing to a non-removable,
18 non-volatile magnetic media (not shown and typically called a "hard drive"), a
19 magnetic disk drive 230 for reading from and writing to a removable, non-volatile
20 magnetic disk 232 (e.g., a "floppy disk"), and an optical disk drive 234 for reading
21 from or writing to a removable, non-volatile optical disk 236 such as a CD-ROM,
22 DVD-ROM or other optical media. The hard disk drive 228, magnetic disk drive
23 230, and optical disk drive 234 are each connected to bus 206 by one or more
24 interfaces 226.

1 The drives and their associated computer-readable media provide
2 nonvolatile storage of computer readable instructions, data structures, program
3 modules, and other data for computer 200. Although the exemplary environment
4 described herein employs a hard disk 228, a removable magnetic disk 232 and a
5 removable optical disk 236, it should be appreciated by those skilled in the art that
6 other types of computer readable media which can store data that is accessible by a
7 computer, such as magnetic cassettes, flash memory cards, digital video disks,
8 random access memories (RAMs), read only memories (ROM), and the like, may
9 also be used in the exemplary operating environment.

10 A number of program modules may be stored on the hard disk 228,
11 magnetic disk 232, optical disk 236, ROM 208, or RAM 210, including, by way of
12 example, and not limitation, an operating system 214, one or more application
13 programs 216 (e.g., multimedia application program 224), other program modules
14 218, and program data 220. In accordance with the illustrated example
15 embodiment of Fig. 2, operating system 214 includes an application program
16 interface embodied as a render engine 222. As will be developed more fully
17 below, render engine 222 is exposed to higher-level applications (e.g., 216) to
18 automatically assemble filter graphs in support of user-defined development
19 projects, e.g., media processing projects. Unlike conventional media processing
20 systems, however, render engine 222 utilizes a scalable, dynamically
21 reconfigurable matrix switch to reduce filter graph complexity, thereby reducing
22 the computational and memory resources required to complete a development
23 project. Various aspects of the innovative media processing system represented by
24 a computer 200 implementing the innovative render engine 222 will be developed
25 further, below.

Continuing with Fig. 2, a user may enter commands and information into computer 200 through input devices such as keyboard 238 and pointing device 240 (such as a "mouse"). Other input devices may include a audio/video input device(s) 253, a microphone, joystick, game pad, satellite dish, serial port, scanner, or the like (not shown). These and other input devices are connected to the processing unit(s) 202 through input interface(s) 242 that is coupled to bus 206, but may be connected by other interface and bus structures, such as a parallel port, game port, or a universal serial bus (USB).

A monitor 256 or other type of display device is also connected to bus 206 via an interface, such as a video adapter 244. In addition to the monitor, personal computers typically include other peripheral output devices (not shown), such as speakers and printers, which may be connected through output peripheral interface 246.

Computer 200 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 250. Remote computer 250 may include many or all of the elements and features described herein relative to computer 200 including, for example, render engine 222 and one or more development applications 216 utilizing the resources of render engine 222.

As shown in Fig. 2, computing system 200 is communicatively coupled to remote devices (e.g., remote computer 250) through a local area network (LAN) 251 and a general wide area network (WAN) 252. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet.

1 When used in a LAN networking environment, the computer 200 is
2 connected to LAN 251 through a suitable network interface or adapter 248. When
3 used in a WAN networking environment, the computer 200 typically includes a
4 modem 254 or other means for establishing communications over the WAN 252.
5 The modem 254, which may be internal or external, may be connected to the
6 system bus 206 via the user input interface 242, or other appropriate mechanism.

7 In a networked environment, program modules depicted relative to the
8 personal computer 200, or portions thereof, may be stored in a remote memory
9 storage device. By way of example, and not limitation, Fig. 2 illustrates remote
10 application programs 216 as residing on a memory device of remote computer
11 250. It will be appreciated that the network connections shown and described are
12 exemplary and other means of establishing a communications link between the
13 computers may be used.

14 Turning next to **Fig. 3**, a block diagram of an example development system
15 architecture is presented, in accordance with one embodiment of the present
16 invention. In accordance with the illustrated example embodiment of Fig. 3,
17 development system 300 is shown comprising one or more application program(s)
18 216 coupled to render engine 222 via an appropriate communications interface
19 302. As used herein, application program(s) 216 are intended to represent any of a
20 wide variety of applications which may benefit from use of render engine 222
21 such as, for example a media processing application 224.

22 The communications interface 302 is intended to represent any of a number
23 of alternate interfaces used by operating systems to expose application program
24 interface(s) to applications. According to one example implementation, interface
25 302 is a component object model (COM) interface, as used by operating systems

1 offered by Microsoft Corporation. As introduced above, COM interface 302
2 provides a means by which the features of the render engine 222, to be described
3 more fully below, are exposed to an application program 216.

4 In accordance with the illustrated example implementation of Fig. 3, render
5 engine 222 is presented comprising source filter(s) 304A-N, transform filter(s)
6 306A-N and render filter 310, coupled together utilizing virtual pins to facilitate a
7 user-defined media processing project. According to one implementation, the
8 filters of system 300 are similar to the filters exposed in conventional media
9 processing systems. According to one implementation, however, filters are not
10 coupled via such interface pins. Rather, alternate implementations are envisioned
11 wherein individual filters (implemented as objects) make calls to other objects,
12 under the control of the render engine 222, for the desired input. Unlike
13 conventional systems, however, render engine 222 exposes a scalable, dynamically
14 reconfigurable matrix switch filter 308, automatically generated and dynamically
15 configured by render engine 222 to reduce the computational and memory
16 resource requirements often associated with development projects. As introduced
17 above, the pins (input and/or output) are application interface(s) designed to
18 communicatively couple other objects (e.g., filters).

19 In accordance with the example implementation of a media processing
20 system, an application communicates with an instance of render engine 222 when
21 the application 216 wants to process streaming media content. Render engine 222
22 selectively invokes and controls an instance of filter graph manager (not shown) to
23 automatically create a filter graph by invoking the appropriate filters (e.g., source,
24 transform and rendering). As introduced above, the communication of media
25 content between filters is achieved by either (1) coupling virtual output pins of one

1 filter to the virtual input pins of requesting filter; or (2) by scheduling object calls
2 between appropriate filters to communicate the requested information. As shown,
3 source filter 304 receives streaming data from the invoking application or an
4 external source (not shown). It is to be appreciated that the streaming data can be
5 obtained from a file on a disk, a network, a satellite feed, an Internet server, a
6 video cassette recorder, or other source of media content. As introduced above,
7 transform filter(s) 306 take the media content and processes it in some manner,
8 before passing it along to render filter 310. As used herein, transform filter(s) 306
9 are intended to represent a wide variety of processing methods or applications that
10 can be performed on media content. In this regard, transform filter(s) 306 may
11 well include a splitter, a decoder, a sizing filter, a transition filter, an effects filter,
12 and the like. The function of each of these filters is described more fully in the
13 Griffiths application, introduced above, and generally incorporated herein by
14 reference. The transition filter, as used herein, is utilized by render engine 222 to
15 transition the rendered output from a first source to a second source. The effect
16 filter is selectively invoked to introduce a particular effect (e.g., fade, wipe, audio
17 distortion, etc.) to a media stream.

18 In accordance with one aspect of the embodiment, to be described more
19 fully below, matrix switch filter 308 selectively passes media content from one or
20 more of a scalable plurality of input(s) to a scalable plurality of output(s).
21 Moreover, matrix switch 308 also supports implementation of a cascaded
22 architecture utilizing feedback paths, i.e., wherein transform filters 306B, 306C,
23 etc. coupled to the output of matrix switch 308 are dynamically coupled to one or
24 more of the scalable plurality of matrix switch input(s). An example of this
25

1 cascaded filter graph architecture is introduced in Fig. 3, and further explained in
2 example implementations, below.

3 Typically, media processed through source, transform and matrix switch
4 filters are ultimately passed to render filter 310, which provides the necessary
5 interface to a hardware device, or other location that accepts the renderer output
6 format, such as a memory or disk file, or a rendering device.

7 **Fig. 4** is a graphical illustration of an example software-enabled matrix
8 switch 308, according to one example embodiment of the present invention. As
9 shown, the matrix switch 308 is comprised of a scalable plurality of input(s) 402
10 and a scalable plurality of output(s) 404, wherein any one or more of the input(s)
11 402 may be iteratively coupled to any one or more of the output(s) 404, based on
12 the content of the matrix switch programming grid 406, automatically generated
13 by render engine 222. According to an alternate implementation introduced
14 above, switch matrix 308 is programmed by render engine 222 to dynamically
15 generate object calls to communicate media content between filters. In addition,
16 according to one implementation, matrix switch 308 includes a plurality of
17 input/output (I/O) buffers 408, as well as means for maintaining source, or media
18 time 410 and/or timeline, or project time 412. It is to be appreciated, however,
19 that in alternate implementations matrix switch 308 does not maintain both source
20 and project times, relying on an upstream filter to convert between these times. As
21 will be developed more fully below, matrix switch 308 dynamically couples one or
22 more of the scalable plurality of inputs 402 to one or more of the scalable plurality
23 of outputs 404 based, at least in part, on the media time 410 and/or the project time
24 412 and the content of matrix switch programming grid 406. In this regard, matrix
25 switch 308 may be characterized as time-aware, supporting such advanced editing

1 features as searching/seeking to a particular point (e.g., media time) in the media
2 content, facilitating an innovative buffering process utilizing I/O buffers 408 to
3 facilitate look-ahead processing of media content, and the like. Thus, it will be
4 appreciated given the discussion to follow that introduction of the matrix switch
5 308 provides a user with an editing flexibility that was heretofore unavailable in a
6 personal computer-based media processing system.

7 As introduced above, the inputs 402 and outputs 404 of matrix switch 308
8 are interfaces which facilitate the time-sensitive routing of data (e.g., media
9 content) in accordance with a user-defined development project. Matrix switch
10 308 has a scalable plurality of inputs 402 and outputs 404, meaning that the
11 number of inputs 402 and outputs 404 are individually generated to satisfy a given
12 editing project. Insofar as each of the inputs/outputs (I/O) has an associated
13 transfer buffer (preferably shared with an adjacent filter) to communicate media
14 content, the scalability of the input/output serves to reduce the overall buffer
15 memory consumed by an editing project. According to one implementation,
16 output 1 is generally reserved as a primary output, e.g., coupled to a rendering
17 filter (not shown).

18 According to one implementation, for each input 402 and output 404,
19 matrix switch 308 attempts to be the allocator, or manager of the buffer associated
20 with the I/O(s) shared with adjacent filters. One reason is to ensure that all of the
21 buffers are of the same size and share common attributes so that a buffer
22 associated with any input 402 may be shared with any output 404, thereby
23 reducing the need to copy memory contents between individual buffers associated
24 with such inputs/outputs. If matrix switch 308 cannot be an allocator for a given
25 output (404), communication from an input (402) to that output is performed using

1 a conventional memory copy operation between the individual buffers associated
2 with the select input/output.

3 As introduced above, the matrix switch programming grid 406 is
4 dynamically generated by render engine 222 based, at least in part, on the user-
5 defined development project. As will be developed below, render engine 222
6 invokes an instance of filter graph manager to assemble a tree structure of an
7 editing project, noting dependencies between source, filters and time to
8 dynamically generate the programming grid 406. A data structure comprising an
9 example programming grid 406 is introduced with reference to Fig. 5, below.

10 Turning briefly to **Fig. 5**, a graphical representation of a data structure
11 comprising an example programming grid 406 is presented, in accordance with
12 one embodiment of the present invention. In accordance with the illustrated
13 example embodiment of Fig. 5, programming grid 406 is depicted as a two-
14 dimensional data structure comprising a column along the y-axis 502 of the grid
15 denoting input pins associated with a content chain (e.g., series of filters to process
16 media content) of the development project. The top row along the x-axis 504 of
17 the data structure denotes project time. With these grid “borders”, the body 506 of
18 the grid 406 is populated with output pin assignments, denoting which input pin is
19 coupled to which output pin during execution of the development project. In this
20 way, render engine 222 dynamically generates and facilitates matrix switch 308.
21 Those skilled in the art will appreciate, however, that data structures of greater or
22 lesser complexity may well be used in support of the programming grid 406
23 without deviating from the spirit and scope of the present invention.

24 Returning to Fig. 4, matrix switch 308 is also depicted with a plurality of
25 input/output buffers 408, shared among all of the input(s)/output(s) (402, 404) to

1 facilitate advanced processing features. That is, while not required to implement
2 the core features of matrix switch 308, I/O buffers 408 facilitate a number of
3 innovative performance enhancing features to improve the performance (or at least
4 the user's perception of performance) of the processing system, thereby providing
5 an improved user experience. According to one implementation, I/O buffers 408
6 are separate from the buffers assigned to each individual input and output pin in
7 support of communication through the switch. According to one implementation,
8 I/O buffers 408 are primarily used to foster look-ahead processing of the project.
9 Assume, for example, that a large portion of the media processing project required
10 only 50% of the available processing power, while some smaller portion required
11 150% of the available processing power. Implementation of the shared I/O buffers
12 408 enable filter graph manager to execute tasks ahead of schedule and buffer this
13 content in the shared I/O buffers 408 until required. Thus, when execution of the
14 filter graph reaches a point where more than 100% of the available processing
15 power is required, the processing system can continue to supply content from the
16 I/O buffers 408, while the system completes execution of the CPU-intensive tasks.
17 If enough shared buffer space is provided, the user should never know that some
18 tasks were not performed in real-time. According to one implementation, shared
19 buffers 408 are dynamically split into two groups by render engine 222, a first
20 group supports the input(s) 402, while a second (often smaller) group is used in
21 support of a primary output (e.g., output pin 1) to facilitate a second, independent
22 output processing thread. The use of an independent output buffers the render
23 engine from processing delays that might occur in upstream and/or downstream
24 filters, as discussed above. It will be appreciated by those skilled in the art that
25

1 such that matrix switch 308 and the foregoing described architecture beneficially
2 suited to support media streaming applications.

3 As introduced above, the filter graph is time-aware in the sense that media
4 (source) time and project execution time are maintained. According to one
5 implementation, matrix switch 308 maintains at least the project clock, while an
6 upstream filter maintains the source time, converting between source and project
7 time for all downstream filters (i.e., including the matrix switch 308). According
8 to one implementation, the frame rate converter filter of a filter graph is
9 responsible for converting source time to project time, and vice versa, i.e.,
10 supporting random seeks, etc. Alternatively, matrix switch 308 utilizes an
11 integrated set of clock(s) to independently maintain project and media times.

12 Having introduced the architectural and operational elements of matrix
13 switch filter 308, **Fig. 6** graphically illustrates an example filter graph
14 implementation incorporating the innovative matrix switch 308. In accordance
15 with the illustrated example embodiment, filter graph 600 is generated by render
16 engine 222 in response to a user defined development project. Unlike the lengthy
17 linear filter graphs typical of convention development systems however, filter
18 graph 600 is shown incorporating a matrix switch filter 308 to recursively route
19 the pre-processed content (e.g., through filters 602, 606, 610, 614 and 618,
20 described more fully below) through a user-defined number of transform filters
21 including, for example, transition filter(s) 620 and effects filter(s) 622. Moreover,
22 as will be developed more fully below, the scalable nature of matrix switch filter
23 308 facilitates such iterative processing for any number of content threads, tracks
24 or compositions.

1 According to one implementation, a matrix switch filter 308 can only
2 process one type of media content, of the same size and at the same frame-rate
3 (video) or modulation type/schema (audio). Thus, Fig. 6 is depicted comprising
4 pre-processing filters with a parser filter 606 to separate, independent content
5 type(s) (e.g., audio content and video content), wherein one of the media types
6 would be processed along a different path including a separate instance of matrix
7 switch 308. Thus, in accordance with the illustrated example embodiment of a
8 media processing system, processing multimedia content including audio and
9 video would utilize two (2) matrix switch filters 308, one dedicated to audio
10 processing (not shown) and one dedicated to video processing. That is not to say,
11 however, that multiple switch filters 308 could not be used (e.g., two each for
12 audio and video) for each content type in alternate implementations. Similarly, it
13 is anticipated that in alternate implementations a matrix switch 308 that accepts
14 multiple media types could well be used without deviating from the spirit and
15 scope of the present invention.

16 In addition filter graph 600 includes a decoder filter 610 to decode the
17 media content. Resize filter 614 is employed when matrix switch 308 is to receive
18 content from multiple sources, ensuring that the size of the received content is the
19 same, regardless of the source. According to one implementation, resize filter 614
20 is selectively employed in video processing paths to adjust the media size of
21 content from one or more sources to a user-defined level. Alternatively, resizer
22 filter 614 adjusts the media size to the largest size provided by any one or more
23 media sources. That is, if, for example, render engine 222 identifies the largest
24 required media size (e.g., 1270x1040 video pixels per frame) and, for any content
25 source not providing content at this size, the content is modified (e.g., stretched,

1 packed, etc.) to fill this size requirement. The frame rate converter (FRC) and
2 pack filter 618, introduced above, ensures that video content from the multiple
3 sources is arriving at the same frame rate, e.g., ten (10) frames per second. As
4 introduced above, the FRC also maintains the distinction between source time and
5 project time.

6 In accordance with one aspect of the present invention, filter graph 600 is
7 depicted utilizing a single, negotiated buffer 604, 608, 612, 616, etc. between
8 adjacent filters. In this regard, render engine 222 reduces the buffer memory
9 requirements in support of a development project.

10 From the point of pre-processing (filters 602, 606, 610, 614, 618), rather
11 than continue a linear filter graph incorporating all of the transition 620 and effect
12 622 filter(s), render engine 222 utilizes a cascade architecture, recursively passing
13 media content through the matrix switch 308 to apply to the transform filter(s)
14 (e.g., 620, 622, etc.) to complete the execution of the development project. It will
15 be appreciated by those skilled in the art that the ability to recursively pass media
16 content through one or more effect and/or transition filters provided by the matrix
17 switch filter 308 greatly reduces the perceived complexity of otherwise large filter
18 graphs, while reducing memory and computational overhead.

19 Turning to **Fig. 7**, a flow chart of an example method for generating a filter
20 graph is presented, in accordance with one aspect of the present invention. The
21 method 700 begins with block 702 wherein render engine 222 receives an
22 indication to generate a filter graph representing a user-defined development
23 project (e.g., a media editing project). According to one example implementation,
24 the indication is received from an application 224 via COM interface(s) 302.
25

1 In block 704, render engine 222 facilitates generation of the editing project,
2 identifying the number and type of media sources selected by the user. In block
3 706, based at least in part on the number and/or type of media sources, filter graph
4 manger 222 exposes source, transform and rendering filter(s) to effect a user
5 defined media processing project, while beginning to establish a programming
6 grid 406 for the matrix switch filter 308.

7 In block 708, reflecting user editing instructions, render engine 222
8 completes the programming grid 406 for matrix switch 308, identifying which
9 inputs 402 are to be coupled to which outputs 404 at particular project times.

10 Based, at least in part, on the programming grid 406 render engine 222
11 generates a matrix switch filter 308 with an appropriate number of input 402 and
12 output 404 pins to effect the project, and assembles the filter graph, block 710.

13 In block 712, to reduce the buffer memory requirements for the processing
14 project, the render engine 222 instructs the filters populating the filter graph to
15 (re)negotiate buffer memory requirements between filters. That is, adjacent filters
16 attempt to negotiate a size and attribute standard so that a single buffer can be
17 utilized to couple each an output pin of one filter to an input pin of a downstream
18 filter. An example implementation of the buffer negotiation process of block 712
19 is presented in greater detail with reference to Fig. 8.

20 Turning briefly to Fig. 8, an example method of negotiating buffer
21 requirements between adjacent filters is presented, in accordance with one
22 example implementation of the present invention. Once the final connection is
23 established to matrix switch 308, matrix switch 308 identifies the maximum buffer
24 requirements for any filter coupled to any of its pins (input 402 and/or output 404),
25 block 802. According to one implementation, the maximum buffer requirements

are defined as the lowest common multiple of buffer alignment requirements, and the maximum of all the pre-fix requirements of the filter buffers.

In block 804, matrix switch 308 selectively removes one or more existing filter connections to adjacent filters. Matrix switch 308 then reconnects all of its pins to adjacent filters using a common buffer size between each of the pins, block 806. In block 808, matrix switch 308 negotiates to be the allocator for all of its pins (402, 404). If the matrix switch 308 cannot, for whatever reason, be the allocator for any of its input pins 402 minimal loss to performance is encountered, as the buffer associated with the input pin will still be compatible with any downstream filter (i.e., coupled to an output pin) and, thus, the buffer can still be passed to the downstream filter without requiring a memory copy operation. If, however, matrix switch 308 cannot be an allocator for one of its output pins 404, media content must then be transferred to at least the downstream filter associated with that output pin using a memory copy operation, block 810.

In block 812, once the matrix switch 308 has re-established its connection to adjacent filters, render engine 222 restores the connection in remaining filters using negotiated buffer requirements emanating from the matrix switch filter 308 buffer negotiations. Once the connections throughout the filter graph have been reconnected, the process continues with block 714 of Fig. 7.

In block 714 (Fig. 7), have re-established the connections between filters, render engine 222 is ready to implement a user's instruction to execute the media processing project.

Example Operation and Implementation(s)

1 The matrix switch described above is quite useful in that it allows multiple
2 inputs to be directed to multiple outputs at any one time. These input can compete
3 for a matrix switch output. The embodiments described below permit these
4 competing inputs to be organized so that the inputs smoothly flow through the
5 matrix switch to provide a desired output. And, while the inventive programming
6 techniques are described in connection with the matrix switch as such is employed
7 in the context of multi-media editing projects, it should be clearly understood that
8 application of the inventive programming techniques and structures should not be
9 so limited only to application in the field of multi-media editing projects or, for
10 that matter, multi-media applications or data streams. Accordingly, the principles
11 about to be discussed can be applied to other fields of endeavor in which multiple
12 inputs can be characterized as competing for a particular output during a common
13 time period.

14 In the multi-media example below, the primary output of the matrix switch
15 is a data stream that defines an editing project that has been created by a user.
16 Recall that this editing project can include multiple different sources that are
17 combined in any number of different ways, and the sources that make up a project
18 can comprise audio sources, video sources, or both. The organization of the inputs
19 and outputs of the matrix switch are made manageable, in the examples described
20 below, by a data structure that permits the matrix switch to be programmed.

21 Fig. 9 shows an overview of a process that takes a user-defined editing
22 project and renders from it a data structure that can be used to program the matrix
23 switch.

24 Specifically, a user-defined editing project is shown generally at 900.
25 Typically, when a user creates an editing project, they can select from a number of

different multimedia clips that they can then assemble into a unique presentation. Each individual clip represents a *source* of digital data or a source stream (e.g., multimedia content). Projects can include one or more sources 902. In defining their project, a user can operate on sources in different ways. For example, video sources can have *transitions* 904 and *effects* 906 applied on them. A transition object is a way to change between two or more sources. As discussed above, a transition essentially receives as input, two or more streams, operates on them in some way, and produces a single output stream. An exemplary transition can comprise, for example, fading from one source to another. An effect object can operate on a single source or on a composite of sources. An effect essentially receives a single input stream, operates on it in some way, and produces a single output stream. An exemplary effect can comprise a black-and-white effect in which a video stream that is configured for presentation in color format is rendered into a video stream that is configured for presentation in black and white format. Unlike conventional effect filters, effect object 906 may well perform multiple effect tasks. That is, in accordance with one implementation, an effect object (e.g., 906) may actually perform multiple tasks on the received input stream, wherein said tasks would require multiple effect filters in a conventional filter graph system.

An exemplary user interface 908 is shown and represents what a user might see when they produce a multimedia project with software executing on a computer. In this example, the user has selected three sources A, B, and C, and has assembled the sources into a project timeline. The project timeline defines when the individual sources are to be rendered, as well as when any transitions and/or effects are to occur.

1 In the discussion that follows, the notion of a *track* is introduced. A track
2 can contain one or more sources or source clips. If a track contains more than one
3 source clip, the source clips cannot overlap. If source clips are to overlap (e.g.
4 fading from one source to another, or having one source obscure another), then
5 multiple tracks are used. A track can thus logically represent a layer on which
6 sequential video is produced. User interface 908 illustrates a project that utilizes
7 three tracks, each of which contains a different source. In this particular project
8 source A will show for a period of time. At a defined time in the presentation,
9 source A is obscured by source B. At some later time, source B transitions to
10 source C.

11 In accordance with the described embodiment, the user-defined editing
12 project 900 is translated into a data structure 910 that represents the project. In the
13 illustrated and described example, this data structure 910 comprises a tree
14 structure. It is to be understood, however, that other data structures could be used.
15 The use of tree structures to represent editing projects is well-known and is not
16 described here in any additional detail. Once the data structure 910 is defined, it is
17 processed to provide a data structure 912 that is utilized to program the matrix
18 switch. In the illustrated and described embodiment, data structure 912 comprises
19 a grid from which the matrix switch can be programmed. It is to be understood
20 and appreciated that other data structures and techniques could, however, be used
21 to program the matrix switch without departing from the spirit and scope of the
22 claimed subject matter.

23 The processing that takes place to define data structures 910 and 912 can
24 take place using any suitable hardware, software, firmware, or combination
25 thereof. In the examples set forth below, the processing takes place utilizing

1 software in the form of a video editing software package that is executable on a
2 general purpose computer.

4 Example Project

5 For purposes of explanation, consider Fig. 10 which shows project 908
6 from Fig. 9 in a little additional detail. Here, a time line containing numbers 0-16
7 is provided adjacent the project to indicate when particular sources are to be seen
8 and when transitions and effects (when present) are to occur. In the examples in
9 this document, the following convention exists with respect to projects, such as
10 project 908. A priority exists for video portions of the project such that as one
11 proceeds from top to bottom, the priority increases. Thus, in the Fig. 10 example,
12 source A has the lowest priority followed by source B and source C. Thus, if there
13 is an overlap between higher and lower priority sources, the higher priority source
14 will prevail. For example, source B will obscure source A from between $t = 4-8$.

15 In this example, the following can be ascertained from the project 908 and
16 time line: from time $t=0-4$ source A should be routed to the matrix switch's
17 primary output; from $t=4-12$ source B should be routed to the matrix switch's
18 primary output; from $t=12-14$ there should be a transition between source B and
19 source C which should be routed to the matrix switch's primary output; and from
20 $t=14-16$ source C should be routed to the matrix switch's primary output. Thus,
21 relative to the matrix switch, each of the sources and the transition can be
22 characterized by where it is to be routed at any given time. Consider, for example,
23 the table just below:

Object	Routing for a given time
C	$t = 0-12$ (nowhere); $t = 12-14$ (transition); $t = 14-16$ (primary output)
B	$t = 0-4$ (nowhere); $t = 4-12$ (primary output); $t = 12-14$ (transition); $t = 14-16$ (nowhere)
A	$t = 0-4$ (primary output); $t = 4-16$ (nowhere)
Transition	$t = 0-12$ (nowhere); $t = 12-14$ (primary output); $t = 14-16$ (nowhere)

Fig. 11 shows an exemplary matrix switch 1100 that can be utilized in the presentation of the user's project. Matrix switch 1100 comprises multiple inputs and multiple outputs. Recall that a characteristic of the matrix switch 1100 is that any of the inputs can be routed to any of the outputs at any given time. A transition element 1102 is provided and represents the transition that is to occur between sources B and C. Notice that the matrix switch includes four inputs numbered 0-3 and three outputs numbered 0-2. Inputs 0-2 correspond respectively to sources A-C, while input 3 corresponds to the output of the transition element 1102. Output 0 corresponds to the switch's primary output, while outputs 1 and 2 are routed to the transition element 1102.

The information that is contained in the table above is the information that is utilized to program the matrix switch. The discussion presented below describes but one implementation in which the information contained in the above table can be derived from the user's project time line.

Recall that as a user edits or creates a project, software that comprises a part of their editing software builds a data structure that represents the project. In the Fig. 9 overview, this was data structure 910. In addition to building the data structure that represents the editing project, the software also builds and configures

1 a matrix switch that is to be used to define the output stream that embodies the
2 project. Building and configuring the matrix switch can include building the
3 appropriate graphs (e.g., a collection of software objects, or filters) that are
4 associated with each of the sources and associating those graphs with the correct
5 inputs of the matrix switch. In addition, building and configuring the matrix
6 switch can also include obtaining and incorporating additional appropriate filters
7 with the matrix switch, e.g. filters for transitions, effects, and mixing (for audio
8 streams). This will become more apparent below.

9 Fig. 12 shows a graphic representation of an exemplary data structure 1200
10 that represents the project of Fig. 10. Here, the data structure comprises a
11 traditional hierarchical tree structure. Any suitable data structure can, however, be
12 utilized. The top node 1202 constitutes a *group* node. A *group* encapsulates a type
13 of media. For example, in the present example the media type comprises video.
14 Another media type is audio. The group node can have child nodes that are either
15 tracks or composites. In the present example, three track nodes 1204, 1206, and
16 1208 are shown. Recall that each track can have one or more sources. If a track
17 comprises more than one source, the sources cannot overlap. Here, all of the
18 sources (A, B, and C) overlap. Hence, three different tracks are utilized for the
19 sources. In terms of priority, the lowest priority source is placed into the tree
20 furthest from the left at 1204a. The other sources are similarly placed. Notice that
21 source C (1208a) has a transition 1210 associated with it. A transition object, in
22 this example, defines a two-input/one output operation. When applied to a track
23 or a composition (discussed below in more detail), the transition object will
24 operate between the track to which it has been applied, and any objects that are
25 beneath it in priority and at the same level in the tree. A “tree level” has a

1 common depth within the tree and belongs to the same parent. Accordingly, in
2 this example, the transition 1210 will operate on a source to the left of the track on
3 which source C resides, and beneath it in priority, i.e. source B. If the transition is
4 applied to any object that has nothing beneath it in the tree, it will transition from
5 blackness (and/or silence if audio is included).

6 Once a data structure representing the project has been built, in this case a
7 hierarchical tree structure, a rendering engine processes the data structure to
8 provide another data structure that is utilized to program the matrix switch. In the
9 Fig. 9 example, this additional data structure is represented at 912. It will be
10 appreciated and understood that the nodes of tree 1200 can include so-called meta
11 information such as a name, ID, and a time value that represents when that
12 particular node's object desires to be routed to the output, e.g. node 1204a would
13 include an identifier for the node associating it with source A, as well as a time
14 value that indicates that source A desires to be routed to the output from time $t = 0$ -
15 8. This meta information is utilized to build the data structure that is, in turn,
16 utilized to program the matrix switch.

17 In the example about to be described below, a specific data structure in the
18 form of a grid is utilized. In addition, certain specifics are described with respect
19 to how the grid is processed so that the matrix switch can be programmed. It is to
20 be understood that the specific described approach is for exemplary purposes only
21 and is not intended to limit application of the claims. Rather, the specific approach
22 constitutes but one way of implementing broader conceptual notions embodied by
23 the inventive subject matter.

24 Figs. 13-18 represent a process through which the inventive grid is built. In
25 the grid about to be described, the x axis represents time, and the y axis represents

1 layers in terms of priority that go from lowest (at the top of the grid) to highest (at
2 the bottom of the grid). Every row in the grid represents the video layer.
3 Additionally, entries made within the grid represent output pins of the matrix
4 switch. This will become apparent below.

5 The way that the grid is built in this example is that the rendering engine
6 does a traversal operation on the tree 1200. In this particular example, the
7 traversal operation is known as a “depth-first, left-to-right” traversal. This
8 operation will layerize the nodes so that the leftmost track or source has the lowest
9 priority and so on. Doing the above-mentioned traversal on tree 1200 (Fig. 12),
10 the first node encountered is node 1204 which is associated with source A. This is
11 the lowest priority track or source. A first row is defined for the grid and is
12 associated with source A. After the first grid row is defined, a grid entry is made
13 and represents the time period for which source A desires to be routed to the
14 matrix switch’s primary output.

15 Fig. 13 shows the state of a grid 1300 after this first processing step.
16 Notice that from time $t = 0-8$, a “0” has been placed in the grid. The “0”
17 represents the output pin of the matrix switch—in this case the primary output.
18 Next, the traversal encounters node 1206 (Fig. 12) which is associated with source
19 B. A second row is thus defined for the grid and is associated with source B.
20 After the second grid row is defined, a grid entry is made and represents the time
21 period for which source B desires to be routed to the matrix switch’s primary
22 output.

23 Fig. 14 shows the state of grid 1300 after this second processing step.
24 Notice that from time $t = 4-14$, a “0” has been placed in the grid. Notice at this
25 point that something interesting has occurred which will be resolved below. Each

1 of the layers has a common period of time (i.e. $t = 4-8$) for which it desires to be
2 routed to the matrix switch's primary output. However, because of the nature of
3 the matrix switch, only one input can be routed to the primary output at a time.
4 Next, the traversal encounters node 1208 (Fig. 12) which is associated with source
5 C. In this particular processing example, a rule is defined that sources on tracks
6 are processed before transitions on the tracks are processed because transitions
7 operate on two objects that are beneath them. A third row is thus defined for the
8 grid and is associated with source C. After the third row is defined, a grid entry is
9 made and represents the time period for which source C desires to be routed to the
10 matrix switch's primary output.

11 Fig. 15 shows the state of grid 1300 after this third processing step. Notice
12 that from time $t = 12-16$, a "0" has been placed in the grid. Next, the traversal
13 encounters node 1210 (Fig. 12) which corresponds to the transition. Thus, a fourth
14 row is defined in the grid and is associated with the transition. After the fourth
15 row is defined, a grid entry is made and represents the time period for which the
16 transition desires to be routed to the matrix switch's primary output.

17 Fig. 16 shows the state of grid 1300 after this fourth processing step.
18 Notice that from time $t = 12-14$, a "0" has been placed in the grid for the transition
19 entry. The transition is a special grid entry. Recall that the transition is
20 programmed to operate on two inputs and provide a single output. Accordingly,
21 starting at the transition entry in the grid and working backward, each of the
22 entries corresponding to the same tree level are examined to ascertain whether
23 they contain entries that indicate that they want to be routed to the output during
24 the same time that the transition is to be routed to the output. If grid entries are
25 found that conflict with the transition's grid entry, the conflicting grid entry is

1 changed to a value to corresponds to an output pin that serves as an input to the
2 transition element 1102 (Fig. 11). This is essentially a redirection operation. In
3 the illustrated grid example, the transition first finds the level that corresponds to
4 source C. This level conflicts with the transition's grid entry for the time period t
5 = 12-14. Thus, for this time period, the grid entry for level C is changed to a
6 switch output that corresponds to an input for the transition element. In this
7 example, a "2" is placed in the grid to signify that for this given time period, this
8 input is routed to output pin 2. Similarly, continuing up the grid, the next level
9 that conflicts with the transition's grid entry is the level that corresponds to source
10 B. Thus, for the conflicting time period, the grid entry for level B is changed to a
11 switch output that corresponds to an input for the transition element. In this
12 example, a "1" is placed in the grid to signify that for this given time period, this
13 input is routed to output pin 1 of the matrix switch.

14 Fig. 17 shows the state of the grid at this point in the processing. Next, a
15 pruning function is implemented which removes any other lower priority entry
16 that is contending for the output with a higher priority entry. In the example, the
17 portion of A from $t=4-8$ gets removed because the higher priority B wants the
18 output for that time.

19 Fig. 18 shows the grid with a cross-hatched area that signifies that portion
20 of A's grid entry that has been removed.

21 At this point, the grid is in a state in which it can be used to program the
22 matrix switch. The left side entries -- A, B, C, and TRANS represent input pin
23 numbers 0, 1, 2, and 3 (as shown) respectively, on the matrix switch shown in Fig.
24 11. The output pin numbers of the matrix switch are designated at 0, 1, and 2 both
25 on the switch in Fig. 11 and within the grid in Fig. 18. As one proceeds through

the grid, starting with source A, the programming of the matrix switch can be ascertained as follows: A is routed to output pin 0 of the matrix switch (the primary output) from $t = 0-4$. From $t = 4-16$, A is not routed to any output pins. From $t = 0-4$, B is not routed to any of the output pins of the matrix switch. From $t = 4-12$, B is routed to the primary output pin 0 of the matrix switch. From $t = 12-14$, B is routed to output pin 1 of the matrix switch. Output pin 1 of the matrix switch corresponds to one of the input pins for the transition element 1102 (Fig. 11). From $t = 14-16$, B is not routed to any of the output pins of the matrix switch. From $t = 0-12$, C is not routed to any of the output pins of the matrix switch. From $t = 12-14$, C is routed to output pin 2 of the matrix switch. Output pin 2 of the matrix switch corresponds to one of the input pins for the transition element 302 (Fig. 3). From $t = 12-14$ the transition element (input pin 3) is routed to output pin 0. From $t = 14-16$, C is routed to output pin 0 of the matrix switch.

As alluded to above, one of the innovative aspects of the matrix switch 308 is its ability to seek to any point in a source, without having to process the intervening content serially through the filter. Rather, matrix switch 308 identifies an appropriate transition point and dumps at least a subset of the intervening content, and continues processing from the sought point in the content.

The ability of the matrix switch 308 to seek to any point in the media content gives rise to certain performance enhancement heretofore unavailable in computer implemented media processing systems. For example, generation of a filter graph by render engine 222 may take into account certain performance characteristics of the media processing system which will execute the user-defined media processing project. In accordance with this example implementation, render engine 222 may access and analyze the system registry of the operating

1 system, for example, to ascertain the performance characteristics of hardware
2 and/or software elements of the computing system implementing the media
3 processing system, and adjust the filter graph construction to improve the
4 perceived performance of the media processing system by the user. Nonetheless,
5 there will always be a chance that a particular instance of a filter graph will not be
6 able to process the media stream fast enough to provide the desired output at the
7 desired time, i.e., processing of the media stream bogs down leading to delays at
8 the rendering filter. In such a case, matrix switch 308 will recognize that it is not
9 receiving media content at the appropriate project time, and may skip certain
10 sections of the project in an effort to "catch-up" and continue the remainder of the
11 project in real time. According to one implementation, when matrix switch 308
12 detects such a lag in processing, it will analyze the degree of the lag and issue a
13 seek command to the source (through the source processing chain) to a future
14 point in the project, where processing continues without processing any further
15 content prior to the sought point.

16 Thus, for the editing project depicted in Fig. 10, the processing described
17 above first builds a data structure (i.e. data structure 1200 in Fig. 12) that
18 represents the project in hierarchical space, and then uses this data structure to
19 define or create another data structure that can be utilized to program the matrix
20 switch.

21 Fig. 19 is a flow diagram that describes steps in a method in accordance
22 with the described embodiment. The method can be implemented in any suitable
23 hardware, software, firmware, or combination thereof. In the illustrated and
24 described embodiment, the method is implemented in software.
25

Step 1900 provides a matrix switch. An exemplary matrix switch is described above. Step 1902 defines a first data structure that represents the editing project. Any suitable data structure can be used, as will be apparent to those of skill in the art. In the illustrated and described embodiment, the data structure comprises a hierarchical tree structure having nodes that can represent tracks (having one or more sources), composites, transitions and effects. Step 1904 processes the first data structure to provide a second data structure that is configured to program the matrix switch. Any suitable data structure can be utilized to implement the second data structure. In the illustrated and described embodiment, a grid structure is utilized. Exemplary processing techniques for processing the first data structure to provide the second data structure are described above. Step 1906 then uses the second data structure to program the matrix switch.

Example Project with a Transition and an Effect

Consider project 2000 depicted in Fig. 20. In this project there are three tracks, each of which contains a source, i.e. source A, B and C. This project includes an effect applied on source B and a transition between sources B and C. The times are indicated as shown.

As the user creates their project, a data structure representing the project is built. Fig. 21 shows an exemplary data structure in the form of a hierarchical tree 2100 that represents project 2000. There, the data structure includes three tracks, each of which contains one of the sources. The sources are arranged in the tree structure in the order of their priority, starting with the lowest priority source on the left and proceeding to the right. There is an effect (i.e. "Fx") that is attached to

1 or otherwise associated with source B. Additionally, there is a transition attached
2 to or otherwise associated with source C.

3 In building the grid for project 2000, the following rule is employed for
4 effects. An effect, in this example, is a one-input/one-output object that is applied
5 to one object—in this case source B. When the effect is inserted into the grid, it
6 looks for any one object beneath it in priority that has a desire to be routed to the
7 primary output of the matrix switch at the same time. When it finds a suitable
8 object, it redirects that object's output from the matrix switch's primary output to
9 an output associated with the effect.

10 As an example, consider Fig. 22 and the grid 2200. At this point in the
11 processing of tree 2100, the rendering engine has incorporated entries in the grid
12 corresponding to sources A, B and the effect. It has done so by traversing the tree
13 2100 in the above-described way. In this example, the effect has already looked
14 for an object beneath it in priority that is competing for the primary output of the
15 matrix switch. It found an entry for source B and then redirected B's grid entry to
16 a matrix switch output pin that corresponds to the effect—here output pin 1.

17 As the render engine 222 completes its traversal of tree 2100, it completes
18 the grid. Fig. 23 shows a completed grid 2200. Processing of the grid after that
19 which is indicated in Fig. 22 takes place substantially as described above with
20 respect to the first example. Summarizing, this processing though: after the effect
21 is entered into the grid and processed as described above, the traversal of tree 2100
22 next encounters the node associated with source C. Thus, a row is added in the
23 grid for source C and an entry is made to indicate that source C desires the output
24 from $t = 12-16$. Next, the tree traversal encounters the node associated with the
25 transition. Accordingly, a row is added to the grid for the transition and a grid

entry is made to indicate that the transition desires the output from $t = 12-14$. Now, as described above, the grid is examined to find two entries, lower in priority than the transition and located at the same tree level as the transition, that compete for the primary output of the matrix switch. Here, those entries correspond to the grid entries for the effect and source C that occur from $t=12-14$. These grid entries are thus redirected to output pins of the matrix switch 308 that correspond to the transition—here pins 2 and 3 as indicated. Next, the grid is pruned which, in this example, removes a portion of the grid entry corresponding to source A for $t=4-8$ because of a conflict with the higher-priority entry for source B.

Fig. 24 shows the resultant matrix switch that has been built and configured as the grid was being processed above. At this point, the grid can be used to program the matrix switch. From the grid picture, it is very easy to see how the matrix switch 308 is going to be programmed. Source A will be routed to the matrix switch's primary output (pin 0) from $t = 0-4$; source B will be redirected to output pin 1 (effect) from $t = 4-14$ and the effect on B will be routed to the output pin 0 from $t = 4-12$. From $t = 12-14$, the effect and source C will be routed to output pins corresponding to the transition (pins 2 and 3) and, accordingly, during this time the transition (input pin 4) will be routed to the primary output (output pin 0) of the matrix switch. From $t = 14-16$, source C will be routed to the primary output of the matrix switch.

It will be appreciated that as the software, in this case the render engine 222, traverses the tree structure that represents a project, it also builds the appropriate graphs and adds the appropriate filters and graphs to the matrix switch. Thus, for example, as the render engine 222 encounters a tree node associated with

1 source A, in addition to adding an entry to the appropriate grid, the software builds
2 the appropriate graphs (i.e. collection of linked filters), and associates those filters
3 with an input of the matrix switch. Similarly, when the render engine 222
4 encounters an effect node in the tree, the software obtains an effect object or filter
5 and associates it with the appropriate output of the matrix switch. Thus, in the
6 above examples, traversal of the tree structure representing the project also enables
7 the software to construct the appropriate graphs and obtain the appropriate objects
8 and associate those items with the appropriate inputs/outputs of the matrix switch
9 308. Upon completion of the tree traversal and processing of the grid, an
10 appropriate matrix switch has been constructed, and the programming (i.e. timing)
11 of inputs to outputs for the matrix switch has been completed.

12 13 **Treatment of “blanks” in a Project**

14 There may be instances in a project when a user leaves a blank in the
15 project time line. During this blank period, no video or audio is scheduled for
16 play.

17 Fig. 25 shows a project that has such a blank incorporated therein. If there
18 is such a blank left in a project, the software is configured to obtain a “black”
19 source and associate the source with the matrix switch at the appropriate input pin.
20 The grid is then configured when it is built to route the black source to the output
21 at the appropriate times and fade from the black (and silent) source to the next
22 source at the appropriate times. The black source can also be used if there is a
23 transition placed on a source for which there is no additional source from which to
24 transition.

Audio Mixing

In the examples discussed above, sources comprising video streams were discussed. In those examples, at any one time, only two video streams were combined into one video stream. However, each project can, and usually does contain an audio component. Alternately, a project can contain only an audio component. The audio component can typically comprise a number of different audio streams that are combined. The discussion below sets forth but one way of processing and combining audio streams.

In the illustrated example, there is no limit on the number of audio streams that can be combined at any one time.

Suppose, for example, there is an audio project that comprises 5 tracks, A-E. Fig. 26 shows an exemplary project. The shaded portions of each track represent the time during which the track is not playing. So, for example, at $t=0-4$, tracks B, D, and E are mixed together and will play. From $t = 4-10$, tracks A-E are mixed together and will play, and the like.

Fig. 27 shows the grid for this project at 2700. Since we are dealing with this composition now, all of the effects and transitions including the audio mixing are only allowed to affect things in this composition. Thus, there is the concept of a boundary 2702 that prevents any actions or operations in this composition from affecting any other grid entries. Note that there are other entries in the grid and that the presently-illustrated entries represent only those portions of the project that relate to the audio mixing function.

Grid 2700 is essentially set up in a manner similar to that described above with respect to the video projects. That is, for each track, a row is added to the grid and a grid entry is made for the time period during which the source on that

1 track desires to be routed to the primary output of the matrix switch. In the
2 present example, grid entries are made for sources A-E. Next, in the same way
3 that a transition or effect was allocated a row in the grid, a “mix” element is
4 allocated a row in the grid as shown and a grid entry is made to indicate that the
5 mix element desires to be routed to the primary output of the matrix switch for a
6 period of time during which two or more sources compete for the matrix switch’s
7 primary output. Note that in this embodiment, allocation of a grid row for the mix
8 element can be implied. Specifically, whereas in the case of a video project,
9 overlapping sources simply result in playing the higher priority source (unless the
10 user defines a transition between them), in the audio realm, overlapping sources
11 are treated as an implicit request to mix them. Thus, the mix element is allocated a
12 grid row any time there are two or more overlapping sources.

13 Once the mix element is allocated into the grid, the grid is processed to
14 redirect any conflicting source entries to matrix switch output pins that correspond
15 to the mix element. In the above case, redirection of the grid entries starts with pin
16 3 and proceeds through to pin 7. The corresponding matrix switch is shown in
17 Fig. 28. Notice that all of the sources are now redirected through the mix element
18 which is a multi-input/one output element. The mix element’s output is fed back
19 around and becomes input pin 15 of the matrix switch. All of the programming of
20 the matrix switch is now reflected in the grid 2700. Specifically, for the indicated
21 time period in the grid, each of the sources is routed to the mix element which, in
22 turn, mixes the appropriate audio streams and presents them to the primary output
23 pin 0 of the matrix switch.

24 Compositions

1 There are situations that can arise when building an editing project where it
2 would be desirable to apply an effect or a transition on just a subset of a particular
3 project or track. Yet, there is no practicable way to incorporate the desired effect
4 or transition. In the past, attempts to provide added flexibility for editing projects
5 have been made in the form of so called “bounce tracks”, as will be appreciated
6 and understood by those of skill in the art. The use of bounce tracks essentially
7 involves processing various video layers (i.e. tracks), writing or moving the
8 processed layers or tracks to another location, and retrieving the processed layers
9 when later needed for additional processing with other layers or tracks. This type
10 of processing can be slow and inefficient.

11 To provide added flexibility and efficiency for multi-media editing projects,
12 the notion of a *composite or composition* is introduced. A composite or
13 composition can be considered as a representation of an editing project as a single
14 track. Recall that editing projects can have one or more tracks, and each track can
15 be associated with one or more sources that can have effects applied on them or
16 transitions between them. In addition, compositions can be nested inside one
17 another.

19 Example Project with Composite

20 Consider, for example, Fig. 29 which illustrates an exemplary project 2900
21 having a composition 2902. In this example, composition 2902 comprises sources
22 B and C and a transition between B and C that occurs between $t = 12-14$. This
23 composition is treated as an individual track or layer. Project 2900 also includes a
24 source A, and a transition between source A and composition 2902 at $t = 4-8$. It
25 will be appreciated that compositions can be much more complicated than the

1 illustrated composition, which is provided for exemplary purposes only.
2 Compositions are useful because they allow the grouping of a particular set of
3 operations on one or more tracks. The operation set is performed on the grouping,
4 and does not affect tracks that are not within the grouping. To draw an analogy, a
5 composition is similar in principle to a mathematical parenthesis. Those
6 operations that appear within the parenthesis are carried out in conjunction with
7 those operations that are intended to operate of the subject matter of the
8 parenthesis. The operations within the parenthesis do not affect tracks that do not
9 appear within the parenthesis.

10 In accordance with the processing that is described above in connection
11 with Fig. 19, a first data structure is defined that represents the editing project.
12 Fig. 30 shows an exemplary data structure 3000 in the form of a hierarchical tree
13 structure. In this example, group node 3002 includes two children—track node
14 3004 and composite node 3006. Track node 3004 is associated with source A.
15 Composite node 3006 includes two children—track nodes 3008 and 3010 that are
16 respectively associated with sources B (3008a) and C (3010a). A transition T2
17 (3012) is applied on source C and a transition T1 (3014) is applied on composition
18 3006.

19 Next, data structure 3000 is processed to provide a second data structure
20 that is configured to program the matrix switch. Note that as the data structure is
21 being programmed, a matrix switch is being built and configured at the same time.
22 In this example, the second data structure comprises a grid structure that is
23 assembled in much the same way as was described above. There are, however,
24 some differences and, for purposes of understanding, the complete evolution of the
25

1 grid structure is described here. In the discussion that follows, the completed
2 matrix switch is shown in Fig. 38.

3 When the rendering engine initiates the depth-first, left-to-right traversal of
4 data structure 3000, the first node it encounters is track node 3004 which is
5 associated with source A. Thus, a first row of the grid is defined and a grid entry
6 is made that represents the time period for which source A desires to be routed to
7 the matrix switch's primary output pin.

8 Fig. 31 shows the state of a grid 3100 after this first processing step. Next
9 the traversal of data structure 3000 encounters the composite node 3006. The
10 composite node is associated with two tracks—track 3008 and track 3010. Track
11 3008 is associated with source B. Accordingly, a second row of the grid is defined
12 and a grid entry is made that represents the time period for which source B desires
13 to be routed to the matrix switch's primary output pin. Additionally, since B is a
14 member of a composition, meta-information is contained in the grid that indicates
15 that this grid row defines one boundary of the composition. This meta-
16 information is graphically depicted with a bracket that appears to the left of the
17 grid row.

18 Fig. 32 shows the state of grid 3100 after this processing step. Next, the
19 traversal of data structure 3000 encounters node 3010 which is associated with
20 source C. Thus, a third row of the grid is added and a grid entry is made that
21 represents the time period for which source C desires to be routed to the matrix
22 switch's primary output pin.

23 Fig. 33 shows the state of grid 3100 after this processing step. Notice that
24 the bracket designating the composition now encompasses the grid row associated
25 with source C. The traversal next encounters node 3012 which is the node

1 associated with the *second* transition T2. Thus, as in the above example, a grid
2 row is added for the transition and a grid entry is made that represents the time
3 period for which the transition desires to be routed to the matrix switch's primary
4 output pin.

5 Fig. 34 shows the state of grid 3100 after this processing step. Notice that
6 the bracket designating the composition is now completed and encompasses grid
7 row entries that correspond to sources B and C and the transition between them.
8 Recall from the examples above that a transition, in this example, is programmed
9 to operate on two inputs and provide a single output. In this instance, and because
10 the transition occurs within a composition, the transition is constrained by a rule
11 that does not allow it to operate on any elements outside of the composition.
12 Thus, starting at the transition entry and working backward through the grid,
13 entries at the same tree level and within the composition (as designated by the
14 bracket) are examined to ascertain whether they contain entries that indicate that
15 they want to be routed to the output during the same time that the transition is to
16 be routed to the output. Here, both of the entries for sources B and C have
17 portions that conflict with the transition's entry. Accordingly, those portions of
18 the grid entries for sources B and C are redirected or changed to correspond to
19 output pins that are associated with a transition element that corresponds to
20 transition T2.

21 Fig. 35 shows the state of grid 3100 after this processing step. The
22 traversal next encounters node 3014 which is the node that is associated with the
23 transition that occurs between source A and composition 2902 (Fig. 29).
24 Processing of this transition is similar to processing of the transition immediately
25 above except for the fact that the transition does not occur within the composition.

1 Because the transition occurs between the composition and another source, one of
2 the inputs for the transition will be the composition, and one of the inputs will be
3 source A (which is outside of the composition). Thus, a grid row is added for this
4 transition and a grid entry is made that represents the time period for which the
5 transition desires to be routed to the matrix switch's primary output pin.

6 Fig. 36 shows the state of grid 3100 after this processing step. At this point
7 then, the grid is examined for entries that conflict with the entry for transition T1.
8 One conflicting grid entry is found for the row that corresponds to source B (inside
9 the composition) and one that corresponds to source A (outside the composition).
10 Accordingly, those portions of the grid row that conflict with transition T1 are
11 changed or redirected to have values that are associated with output pins of the
12 matrix switch that are themselves associated with a transition element T1. In this
13 example, redirection causes an entry of "3" and "4" to be inserted as shown.

14 Fig. 37 shows the state of grid 3100 after this processing step. If necessary,
15 a pruning operation would further ensure that the grid has no competing entries for
16 the primary output of the matrix switch. The associated input pin numbers of the
17 matrix switch are shown to the left of grid 3100.

18 Fig. 38 shows a suitably configured matrix switch that has been build in
19 accordance with the processing described above. Recall that, as data structure
20 3000 (Fig. 30) is processed by the rendering engine, a matrix switch is built and
21 configured in parallel with the building and processing of the grid structure that is
22 utilized to program the matrix switch. From the matrix switch and grid 3100 of
23 Fig. 37, the programming of the switch can be easily ascertained.

Fig. 38a shows an exemplary data structure that represents a project that illustrates the usefulness of composites. In this example, the project can mathematically be represented as follows:

$$(Fx\text{-noisy } (A \text{ Tx-Blend } B)) \text{ Tx-Blend } C$$

Here, an effect (noisy) is applied to A blended with B, the result of which is applied to a blend with C. The composite in this example allows the grouping of the things beneath it so that the effect (noisy), when it is applied, is applied to everything that is beneath it. Notice that without the composite node, there is no node where an effect can be applied that will affect (A Tx-Blend B). Hence, in this example, operations that appear within the parenthesis are carried out on tracks that appear within the parenthesis. Those operations do not affect tracks that are not within the parenthesis.

Fig. 39 is a flow diagram that described steps in a method in accordance with one embodiment. The method can be implemented in any suitable hardware, software, firmware, or combination thereof. In the presently-described example, the method is implemented in software.

Step 3900 defines a multimedia editing project that includes at least one composite. The composite represents multiple tracks as a single track for purposes of the processing described just below. It is important to note that, in the processing described just below, and because of the use of composites, the extra processing that is required by bounce tracks is avoided (i.e. operating on two tracks, moving the operation result to another location, and retrieving the operation result when later needed). This reduces the processing time that is

1 required to render a multi-media project. Step 3902 defines a first data structure
2 that represents the editing project. Any suitable data structure can be utilized. In
3 the present example, a data structure in the form of a hierarchical tree is utilized.
4 An exemplary tree is shown in Fig. 30. Step 3904 processes the first data structure
5 to provide a second data structure that is configured to program a matrix switch.
6 In the illustrated example, the second data structure comprises a grid structure.
7 Exemplary processing is described in the context of Figs. 30-37. Step 3906 then
8 programs the matrix switch using the second data structure.

9 10 **Dynamic Graph Building**

11 Having introduced the various architectural and implementation elements
12 of the present invention, above, attention is now drawn to Figs. 40-44, wherein
13 another aspect of the illustrated embodiment is presented. As introduced above,
14 each matrix switch filter 308 is time aware. That is, according to one
15 implementation, matrix switch 308 maintains one or both project time and media
16 source time information. This enables the matrix switch 308 to, among other
17 things, throttle delivery of media content to the matrix switch 308.

18 As an extension of this capability, in accordance with one aspect of the
19 present invention, render engine 222, via matrix switch filter 308, dynamically
20 builds a filter graph representation of a project during execution of the filter graph.
21 That is, render engine 222 based, at least in part by the control performed by
22 matrix switch 308, dynamically loads filter graph chains as they are needed.
23 Further, render engine 222 may well discard, or cache processing chains when
24 they are no longer required to support execution of the processing project. To
25 illustrate the benefits afforded by dynamic graph building, assume, for example,

1 that an editing project included over 100 sources, yet only three (3) of them were
2 ever required at any given time to support execution of the filter graph. Those
3 skilled in the art will appreciate that loading three sources will be executed much
4 fast than 100 sources, thereby permitting execution of the filter graph to
5 commence much more rapidly than conventional filter graph implementations.
6 Further, the memory and processing resources required to support three (3) sources
7 will generally be less than those required to support 100 sources. Thus, those
8 skilled in the art will appreciate that the dynamic graph building properties of the
9 present invention reduce the computational and memory requirements placed on
10 the host system (e.g., computing system 200).

11 **Fig. 40** is a flow chart of an example method for processing media content,
12 in accordance with one embodiment of the present invention. More particularly,
13 Fig. 40 illustrates an example method wherein render engine 222 dynamically
14 generates and manages a filter graph to reduce the computational and/or memory
15 requirements placed on a host system. As shown, method 4000 begins with block
16 4002, wherein render engine 222 receives an indication to generate a development
17 project. According to one implementation, as discussed above, render engine 222
18 receives the indication from a higher-level application 216, e.g., media processing
19 application 224, to assist a user in generating a processing project (e.g., a media
20 processing project).

21 In block 4004, render engine 222 identifies the number and nature of the
22 media sources within the user-defined processing project, in preparation for
23 generating a filter graph representation of the processing project. As introduced
24 above, for each of the identified sources, render engine 222 determines the
25 necessary transform filters 306 required to pre-process the source (i.e., the source

processing chain), preparing the chain for presentation to the matrix switch filter 308 and one or more transition/effect filters 306. Unlike conventional implementations which would proceed to generate the entire filter graph in preparation for execution of the processing project, render engine 222 generates a list of sources and when they are required in the filter graph. An example of a data structure comprising a list of processing chains is presented with reference to Fig. 41.

Turning briefly to **Fig. 41**, a graphical illustration of an example data structure comprising a processing chain execution list is presented. As shown, the chain execution list 4100 is comprised of a number of information fields, e.g., 4102-4110 which detail, in part, which chains are required at a particular time in project execution. In accordance with the illustrated example embodiment of Fig. 41, chain execution list 4100 is depicted comprising a chain identifier field 4102, a source identifier field 4104, a project time field 4106, a source time field 4108, and a dependencies field 4110.

Upon identifying a project source and the associated filters required for pre-processing the source (i.e., the source chain), render engine 222 assigns the chain an identifier which uniquely identifies the source chain within the context of the filter graph. Accordingly, the chain execution list 4100 includes a field 4102 which maintains a list of the chains utilized in the associated project. Within the context of the filter graph, the chain identifier corresponds to the source and the associated pre-processing filters.

The source identifier field 4104 contains information denoting the project source associated with a particular chain identifier. In this regard, the source

1 identifier field 4104 may well contain a file name, a file handle, or any other
2 suitable source identifier.

3 The project time field 4106 denotes at what point during project execution
4 the source chain is required. The source time field 4108 denotes what portion of
5 the source file is required to support execution of the processing project. It should
6 be appreciated that a user may well utilize the whole source file or any part
7 thereof, as defined by the processing project.

8 The dependencies field 4110 denotes whether the associated chain is
9 dependent upon any other chain. As will be described in greater detail below,
10 multiple chains may rely on a common source and/or a subset of another source
11 chain. In certain implementations, it would not be advantageous to unload source
12 chains prior to their execution and/or the execution of chains dependent thereon.
13 Accordingly, render engine 222 maintains a list of such dependencies within the
14 chain execution list 4100. It is to be appreciated, however, that certain
15 circumstances may arise where it is necessary to unload a chain prior to or during
16 execution, or prior to execution of an otherwise dependent chain. One such
17 example is where the processing project utilizes a hierarchical structure, wherein
18 individual chains are assigned a priority level. An implementation is
19 contemplated, for example, wherein the priority of a particular chain is
20 dynamically managed by a matrix switch filter 308 within a filter graph based, at
21 least in part, on how soon the chain is required to support the uninterrupted
22 execution of the processing project, i.e., chains which are required more urgently
23 are assigned a higher priority and, as a result, are processed at the disadvantage of
24 other, lower priority chains. In the extreme, lower priority chains are unloaded to
25 enable loading of a higher priority chain. It is to be appreciated that, although

1 depicted as a two-dimensional data structure, chain execution lists of greater or
2 lesser complexity may well be substituted without deviating from the spirit and
3 scope of the present invention.

4 Returning to Fig 40 and, in particular, block 4006, render engine 222
5 dynamically generates and manages a filter graph representation of the processing
6 project invoking only those chains associated with sources that are necessary to
7 support the current and/or impending execution of the filter graph. It is to be
8 appreciated that by not opening each of the chains of a processing project, render
9 engine 222 reduces the amount of memory required to build the filter graph,
10 thereby reducing the amount of memory required to complete execution of the
11 project, i.e., recall the example where the entire graph utilized 100 sources, but
12 only required three (3) at any given time. An example method of dynamically
13 generating and managing a filter graph is presented with reference to the flow
14 chart illustrated in Fig. 42.

15 Turning to **Fig. 42**, an example method for dynamically generating and
16 managing a filter graph is presented, in accordance with one embodiment of the
17 present invention. In accordance with the illustrated example implementation of
18 Fig. 42, method 4006 commences with block 4202 wherein render engine 222
19 determines which chains are required to fulfill execution of the development
20 project for the next M seconds. According to this example implementation, M
21 must be greater than the minimum time it takes to completely load the next chain.
22 In accordance with the illustrated example implementation, wherein matrix switch
23 filter 308 controls the pace of project execution, matrix switch filter 308 provides
24 an indication to render engine 222 of what chains are required in block 4202.

1 According to one implementation, M is dynamically generated based on a
2 number of factors including, but not limited to, processing speed, available
3 memory, the complexity of the development project, the number and type of the
4 source chains, and the like. In certain implementations, processing system 300
5 maintains a performance history (not shown), and dynamically modifies the
6 processing threshold M based on past performance. According to one
7 implementation, M is stochastically set to ten (10) seconds. Accordingly, render
8 engine 222 maintains only the chains currently required to support the next ten
9 seconds of execution. It is important to note that project execution does not
10 necessarily correlate to rendering of the composite generated by the filter graph.
11 That is, in certain implementations, execution of the filter graph is performed as
12 fast as possible, utilizing the shared memory resources of the matrix switch filter
13 308 to buffer the composite until the rendering chain consumes the composite.

14 In block 4204, render engine 222 determines whether a threshold of loaded
15 chains (e.g., a maximum chain-count) (T) has been exceeded. In certain
16 implementations, the number of loaded chains will be limited due to memory
17 limitations. According to one example implementation, setting T equal to one (1)
18 is popular in that it requires the render engine to analyze the filter graph for chains
19 that are no longer required (e.g., exhausted chains) whenever a new chain is
20 considered for loading. According to one example implementation, the maximum
21 number of loaded chains supported by render engine 222 is seventy (70).
22 Accordingly, once render engine 222 has identified the chains required (block
23 4202), a determination is made of whether there is space in which to load them
24 into the filter graph (block 4204).

1 If the chain count threshold (T) has not yet been reached, render engine 222
2 loads the identified chains, block 4206. Matrix switch filter 308 will initiate
3 execution of the newly loaded chains to fulfill the execution requirements of the
4 development project.

5 If, in block 4204 the chain-count threshold (T) has been reached, render
6 engine 222 determines whether one or more chains may be unloaded from the
7 filter graph. Thus, in block 4206, render engine 222 identifies any currently loaded
8 chains that will not be utilized in the next N seconds. Source chains may be
9 accessed multiple times to process multiple portions of an associated source.
10 Thus, in accordance with steps 4202-4206, a source chain may have been loaded
11 to meet an impending execution requirement, and remains loaded to satisfy a
12 subsequent processing task. However, where resources are running short, render
13 engine 222 along with matrix switch filter(s) 308 determine which chains are not
14 required in the next N seconds and, in block 4210, instructs render engine 222 to
15 unload the identified chains. As above, N may well be dynamically derived based
16 on past performance. In accordance with one example implementation, N is thirty
17 (30) seconds. According to one implementation, render engine 222 determines
18 whether the chains will be required for subsequent processing in the current or a
19 future filter graph. If so, the filter chain is removed from the active filter graph by
20 render engine 222 and cached for subsequent re-integration in this or a future filter
21 graph.

22 In block 4212, render engine 222 determines whether unloading of the
23 identified chain(s) in block 4210 has brought the total chain-count below the
24 threshold a cutoff threshold (V). According to one implementation, V is greater
25 than T. This is particularly useful if T has been set to one (1), as described above.

1 If so, processing continues with block 4206 as render engine 222 loads the chains
2 identified in block 4202.

3 If, in block 4212, the chain-count threshold is still exceeded, render engine
4 222 re-analyzes the current filter graph and identifies the lowest priority chains,
5 block 4214. That is, the filter graph may well be comprised of seventy chains, all
6 of which will be required in the next thirty (N) seconds. If, however, the chains
7 identified in block 4202 are needed prior to any of the seventy chains currently
8 loaded in the filter graph, those chains are assigned a lower priority. Processing
9 continues with block 4210 as the lower priority chains are unloaded, as render
10 engine 222 re-analyzes the chain-count, in block 4212. If the filter graph has
11 space available, processing continues with block 4206, else it continues with block
12 4214.

13 **Fig. 43** graphically illustrates an example data structure utilized to manage
14 dynamic graph building, according to one example implementation. In accordance
15 with the illustrated example embodiment of Fig. 43, a filter graph 4300 is
16 presented. Unlike conventional filter graph implementations, wherein all chains
17 4302-4308 would be loaded prior to execution of the development project, filter
18 graph 4300 illustrates the dynamic nature of the present invention. In the
19 illustrated example of Fig. 43, matrix switch filter 308 has identified at least two
20 source chains 4302, 4304 which are required in the next M seconds to support the
21 timely processing of the development project. Such chains are illustrated in Fig.
22 43 with a solid black line to denote that these chains are currently loaded into the
23 filter graph 4300. In accordance with one aspect of the present invention, the
24 development project may well contain additional chains (e.g., 4306 and 4308) that
25 will be required to complete execution of the development project, but which are

not yet required and are, thus, not yet loaded. Such chains are illustrated in Fig. 43 with dotted lines, denoting that they are not currently loaded into the filter graph 4300. By limiting the number of currently loaded chains to a threshold (T) or, alternatively (V), the present implementation reduces the memory requirements necessary to satisfy even the most complex of development projects by unloading chains when they are no longer required, without stifling the user's creativity by artificially limiting the size of the filter graph.

Fig. 44 is an example filter graph denoting chain dependencies. In accordance with the illustrated example of Fig. 44, filter graph 4400 depicts two chains 4402 and 4404, each coupled to an associated matrix switch filter through an innovative parser object described more fully in a co-pending patent application entitled *A System and Related Methods for Reducing the Instances of Source Files in a Filter Graph*, filed on December 6, 2000 by the inventors of the present application, the disclosure of which is hereby incorporated by reference.

Filter graph 4400 is representative of the situation wherein the source is a multimedia file containing both audio and video content, each requiring a dedicated pre-processing chain and a matrix switch 308. It is to be appreciated that in such a situation, where the video chain 4202 and the audio chain 4204 share a single instance of a source filter, a dependent relationship with respect to unloading the source filter chain. In such a circumstance, it might not be desirable for the video matrix switch, for example, to unload chain 4202 upon completion if, for example, the audio matrix switch requires additional content from the source. Thus, a conventional approach to such a situation would be to invoke separate processing chains, one for the audio content and one for the video content, each having a unique instance of the source filter. The conventional approach has the

1 disadvantage of wasting memory by invoking multiple instances of the same
2 source. To alleviate this problem, the render engine 222 identifies such
3 dependencies within the chain execution list and will not seek to unload the
4 source filter until both processing chains no longer require content from the source
5 (e.g., for at least the next N seconds, as introduced above). In this regard, the
6 present invention alleviates the need to construct two complete chains accessing a
7 common source, thereby reducing the memory requirements necessary to support
8 more complex processing projects.

9 Although the invention has been described in language specific to structural
10 features and/or methodological steps, it is to be understood that the invention
11 defined in the appended claims is not necessarily limited to the specific features or
12 steps described. Rather, the specific features and steps are disclosed as preferred
13 forms of implementing the claimed invention.